

CLAIMS

1. A method of forwarding data packets in a switched communications network, the network including a plurality of end systems and switches connected by links, the switches having access ports connected to end systems and network ports connected to other switches, and each end system having a unique physical address, the method comprising the steps of:

each one of the switches maintaining a local directory of mappings only for access ports or end systems attached to the access ports of the one switch;

when a first packet is received on an access port of a first switch, the first switch accessing its local directory to determine whether the first packet can be forwarded.

2. The method of claim 1, wherein the local directory includes mappings of at least one of:

physical addresses of end systems locally attached to the one switch;

alias addresses of end systems locally attached to the one switch;

virtual LAN identifiers (VLAN-IDs) for the access ports or end systems attached to the one switch; and

VLAN policy.

3. The method of claim 1, wherein:

each one of the switches maintains a remote directory of mappings of access ports or end systems attached to other switches, wherein the mappings in the remote directory are limited to active or attempted connectivity resolutions.

4. The method of claim 3, wherein the mappings in the remote directory include at least one of:

physical addresses of the other switches that provided the connectivity resolutions;

physical addresses of end systems attached to the other switches;
alias addresses of end systems attached to the other switches;
VLAN-IDs for access ports or end systems attached to the other switches; and
VLAN policy.

5. The method of claim 3, wherein:

when a mapping is not found at the first switch in either one of the local or remote directories, the first switch forwarding a message to all other switches, the message being a request for resolution from the local directory of another switch for a destination end system.

6. The method of claim 1, wherein each switch exchanges switch link state topology information.

7. The method of claim 6, wherein the first switch uses its switch link state topology information to determine a path from a source switch to a destination switch.

8. The method of claim 7, wherein the switch link state topology information includes switch/port pairs and each switch has a single physical layer address, and wherein each switch/port pair comprises the physical address of a switch and a port instance.

9. The method of claim 8, wherein the first switch sends a connect message with a physical address of source end system and a physical address of destination end system as a connection identifier for mapping connections in all switches on the path.

10. The method of claim 9, wherein the connect message includes an in-order list of switch/port pairs on the path.

11. The method of claim 10, wherein the first packet is not forwarded by the first switch until the first switch receives a connect message response from the last switch on the path.

12. The method of claim 11, wherein the connect message maps connections as it travels through the switches on the path, but the mapped connections remain disabled until receipt of a response sent by the last switch on the path which travels back through each switch on the reverse path to the first switch.

13. The method of claim 7, wherein each one of the switches on the path maintains a "links-in-use" table of all links in the path.

14. The method of claim 13, wherein upon receipt of a predetermined link state topology message regarding one or more of the links in the path, each switch on the path accesses its links-in-use table and unmaps any connection for the one or more links identified in the predetermined link state topology message.

15. The method of claim 14, wherein the predetermined link state topology message indicates that a link has failed.

16. The method of claim 1, wherein the path determined by the first switch includes multiple active paths.

17. The method of claim 16, wherein the multiple paths are for the same source switch and destination switch.

18. The method of claim 1, wherein the first switch attempts to resolve broadcast packets in order to reduce the amount of broadcast traffic on the network.

19. The method of claim 18, wherein the first switch attempts to resolve a broadcast packet by looking inside the packet and decoding the protocol layers and higher-level addressing to determine the physical address of the destination.

20. The method of claim 1, wherein unresolvable packets are sent by the first switch on a virtual spanning tree to all other switches, and the packet is tagged with a VLAN-ID which limits the egress ports on the other switches through which the packet is forwarded.

21. A method of mapping connections in a switched communications network, the method comprising the steps of:

determining a path through one or more switches in the network;
a call-originating switch sending a connect request which maps connections as it travels through each switch on the path, but the connections remain disabled until the last switch on the path returns a connect response which enables each connection as it travels through each switch on the path.

22. The method of claim 21, wherein the connect request includes an in-order list of switch/port pairs on the path.

23. The method of claim 22, wherein the connect request travels hop-by-hop according to the in-order list of switch/port pairs on the path.

24. The method of claim 23, wherein each switch has a single physical address and the switch/port pair comprises the physical address of the switch and a port instance.

25. A method of determining the topology in a switched communications network, the method comprising:

each switch exchanging switch link state topology information so that each switch has its own topology graph of the switched network; and

when a first packet is received on an access port of a first switch, the first switch determines a complete path through the switched network by accessing its topology graph.

26. The method of claim 25, wherein the switch link state topology information includes switch/port pairs, where each switch has a single physical address and each switch/port pair comprises a switch physical address and a port instance.

27. The method of claim 26, wherein each link and switch in the switch link topology is automatically named by its switch/port pair.

28. The method of claim 25, wherein the first switch sends a connect message with the path to all switches on the path, and the connect message includes a physical address of a source end system and a physical address of a destination end system as a connection identifier.

29. The method of claim 28, wherein the connect message includes an in-order list of switch/port pairs on the path.

30. A method of topology exchange in a switched communications network, the method comprising:

each switch exchanging switch link state topology information so that each switch has its own topology graph of the switched network, wherein each switch has a single physical address and the topology exchange information includes switch/port pairs, each switch/port pair comprising a switch physical address and a port instance.

31. A method of rerouting connections in a switched communications network, the method comprising the steps of:

- each access switch on an active call path maintaining a "links-in-use" table of all links in the active path;
- upon receipt of a predefined link state topology message regarding one or more links in the path, each switch on the path accessing its links-in-use directory and unmapping a connection for the one or more links identified in the predefined link state topology message.

32. The method of claim 31, wherein the predefined link state message indicates a failed link.

33. The method of claim 31, wherein the active path connection is originated by a call-originating switch, and when the call-originating switch receives the predefined link state message, it redetermines a new active path connection.

34. A method of forwarding data packets in a switched communications network, the network including a plurality of end systems and switches connected by links, the switches having access ports connected to end systems and network ports connected to other switches, and each end system having a unique physical address, the method comprising the steps of:

- providing a connection-oriented service for forwarding of data packets, including, for each active call, mapping a connection to one or more of the switches;
- each switch maintaining a links-in-use table of each active link included in any active call;
- each switch exchanging switch link state messages concerning the state of the links between network ports; and

when a predefined link state message is received regarding an active link, unmapping any connection having a link in the links-in-use table for the active link identified in the predefined message.

35. A method of forwarding data packets in a switched communications network, the network including a plurality of source and destination end systems and switches connected by links, the switches having access ports connected to end systems and network ports connected to other switches, and each end system having a unique physical address, the method comprising the steps of:

when a first packet is received on an access port of a first switch, the first switch independently determines a call path from a source end system to a destination end system through one or more of the switches based on a unique connection identifier which includes the source and destination end system physical addresses.

36. The method of claim 35, wherein the first switch enables directory resolution, topology exchange, and connection services.

37. The method of claim 35, wherein the first switch independently determines a new call path upon receipt of a predefined switch link state topology message.

38. The method of claim 35, wherein the call path includes multiple paths.

39. The method of claim 38, wherein the multiple paths are for the same source switch and destination switch.

40. A method of request/response messaging in a tree-based network node topology, comprising:

an originating node sending a request message to all other nodes on the tree, wherein during each recursive walking of the tree, a first node sends the request message to adjacent nodes, and waits for a response from each adjacent node, but forwards only one response back toward the originating node, such that the originating node receives only a single response.

41. The method of claim 40, wherein the nodes are switches in a switched communications network, and the resolve message is an interswitch resolve message for resolving a destination end system.

42. A method of determining reachability outside of a switch domain, comprising:

each switch in the domain listening to server and network router advertisements on its access ports; and

the access switch attached to the server or router issuing the advertisements, determining and maintaining a best route from the access switch to the server or router.

43. The method of claim 42, wherein the access switch makes a path determination based on combined router best route metrics and switch best path metrics.

44. A method of forwarding data packets in a switched communications network, the network including a plurality of end systems and switches connected by links, the switches having access ports connected to end systems and network ports connected to other switches, and each end system having a unique physical address, the method comprising the steps of:

each one of the access switches maintaining a local directory of mappings only for access ports or end systems attached to the access ports of the one access switch to one or more virtual LAN identifiers (VLAN-IDs);

when a first packet is received on an access port of a first switch, the first switch accessing its local directory to determine whether the packet can be forwarded.

45. The method of claim 44, wherein each one of the access switches maintains a remote directory of mappings of one or more VLAN-IDs to remote end systems connected to the one access switch, wherein the mappings in the remote directory are limited to active or attempted connectivity resolutions.

46. The method of claim 45, wherein when a mapping is not found in either of the local or remote directories, forwarding a message to all other switches, the message being a request for resolution from the local directory of another switch for a destination end system.

47. A method of forwarding data packets in a switched communications network, comprising:

when a first packet is received on an access port of a first switch, the first switch applying a policy of preserving VLAN boundaries before forwarding the first packet.

48. A method of forwarding data packets in a switched communications network, comprising:

when a first packet is received on an access port of a first switch, the first switch resolving the first packet to a unicast packet from a source to a destination, and the first switch applying a policy which enables forwarding of the packet across VLAN boundaries.

49. A method of forwarding data packets in a switched communications network, comprising:

when a first packet is received on an access port of a first switch, the first switch attempting to resolve the first packet to a unicast destination packet, and if it cannot be so resolved, tagging the first packet with a VLAN identifier and forwarding the packet.

50. A method of forwarding data packets in a switched communications network, comprising:

assigning virtual LAN identifiers (VLAN-IDs) to access ports and end systems attached to access ports; and
wherein an access port may be assigned to more than one VLAN-ID.

51. The method of claim 50, wherein an end system may be assigned more than one VLAN-ID.

52. An apparatus for forwarding data packets in a switched communications network, the network including a plurality of end systems and switches connected by links, the switches having access ports connected to end systems and network ports connected to other switches, and each end system having a unique physical address, the apparatus comprising:

each one of the switches having means for maintaining a local directory of mappings only for access ports or end systems attached to the access ports of the one switch;

when a first packet is received on an access port of a first switch, the first switch having means for accessing its local directory to determine whether the first packet can be forwarded.